ANONYMOUS ATTACKS: Members of the hacker group Anonymous have targeted distributed denial-of-service attacks at companies that have ceased serving WikiLeaks.
PHOTO: VINCENT DIAMANTE

Network Defense Gone Wrong

Some distributed denial-of-service defenses could in fact make a Web site more vulnerable

After WikiLeaks began publishing confidential communications of the U.S. State Department late last year, a spate of incidents put distributed denial-of-service attacks back in the news. These attacks can take many forms, the most straightforward of which is simply to overwhelm the targeted file server with requests. If that server can't keep up with the barrage, legitimate users are effectively shut out.

An attractive defense is to employ a large number of servers at far-flung locations. Such content-delivery networks are common enough, and plugging into one isn't difficult. Akamai Technologies runs the largest such network; it has 73 000 globally distributed servers, which, according to the company, handle 15 to 30 percent of all Web traffic.

With such file-serving clout in your corner, your site should be able to stand up to almost any pummeling, right? That's what Akamai boasted in the wake of the recent attacks. And in most situations, linking up with a content-delivery network is undoubtedly a good defense. But many of the companies doing so probably don't know that if the bad guys are clever enough—and if the good guys are not quite on the ball—using a content-delivery network might actually *increase* a Web site's vulnerability.

That troubling possibility came to light in 2009 at the 14th European Symposium on Research in Computer Security in work reported by Michael Rabinovich, a professor of electrical engineering and computer science at Case Western Reserve University, in Cleveland, and two graduate students then studying under his direction, Sipat Triukose and Zakaria Al-Qudah.

"Content-delivery networks have this intuitively understood claim that they improve the resiliency of Web sites to distributed denial-of-service attacks," says Rabinovich. "If an attacker tries to launch an attack, he will exhaust his resources before the content-delivery network notices a blip." But while Rabinovich's group was studying the performance of content-delivery networks, they stumbled on a subtle weakness. To appreciate the trouble it could cause, you need to understand a little bit about these networks.

When your computer requests something from the Web site of a company that uses a content-delivery network, you are invisibly routed to one of the network's many so-called edge servers. If that edge server does not have the content you are looking for, which might include rather large files, such as video clips or images, it retrieves a copy from the company's server (the origin server) and passes it to you. The edge server also keeps a copy in its cache. Subsequent requests for that content can thus be accommodated without consulting the origin server. This is what protects that server from being swamped with requests.

The basic problem, Rabinovich's team found, is that a bad guy can add what's known as a query string to the URL he is targeting. Query strings are common enough—you often see them at the top of your browser introduced by a question mark. They are used to communicate parameters to the server, such as the particular keywords you are Googling.

The conundrum here is that if a random query string is added to a URL, the content-delivery network's server will typically treat the request as new and pass it on to the origin server. If the origin server is not expecting a query string, it will most likely disregard it and just supply the file normally. That is to say, an attacker can force an edge server to consult the origin server—perhaps to ask for a copy of a large file.

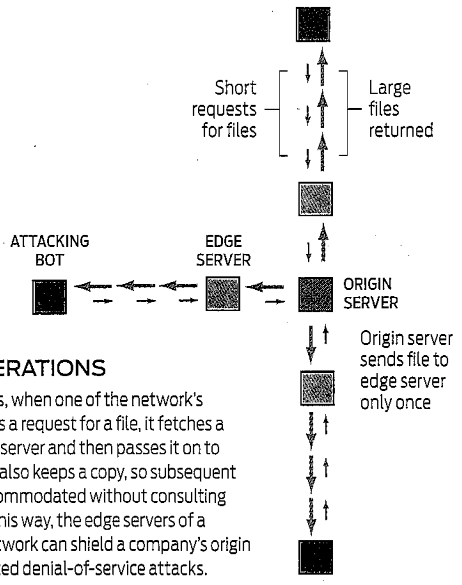

NORMAL OPERATIONS

In normal operations, when one of the network's edge servers receives a request for a file, it fetches a copy from the origin server and then passes it on to the requester. But it also keeps a copy, so subsequent requests can be accommodated without consulting the origin server. In this way, the edge servers of a content-delivery network can shield a company's origin server from distributed denial-of-service attacks.

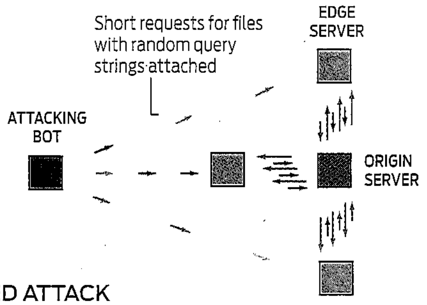

AMPLIFIED ATTACK

A clever attacker could append random query strings to his requests for a file. He could also terminate the connection after making the request to one edge server and then issue more requests to other edge servers. Unless they are configured properly, those edge servers will see each of these requests as being for novel information and so will consult the origin server, which will most likely just ignore the query string and serve a file. In this way the attacker can enlist the edge servers to amplify his attack.

Actually, it's worse than that. The attacker's computer can make such a request and then swiftly terminate the connection. The edge server, however, still contacts the origin server to obtain the requested file, which consumes the origin server's computing resources. Meanwhile, the attacker moves on to make more such requests through the content-delivery network's other edge servers, which, Rabinovich's group also showed, the attacker can reach at will.

Because the attacker can issue these requests with relatively short messages and doesn't have to wait around for files to be returned, little computing power is needed. Meanwhile, the origin server could be inundated with requests from hundreds or even thousands of edge servers asking it to provide large files. In this way, an attacker could effectively enlist the content-delivery network in his assault.

"Before I submitted the paper, I contacted Akamai," says Rabinovich. "As a good Internet citizen, I thought I should let them know before letting the cat out of the bag." Akamai's response, according to Rabinovich, was that it already provides its customers with all they need to guard against such attacks. "Akamai was putting the problem on the content provider," says Rabinovich. "That's sort of like talking the problem away."

Bruce Maggs, vice president for research at Akamai and a professor of computer science at Duke University, in Durham, N.C., sees things differently. He views the vulnerability that Rabinovich pointed out as only a minor worry. Because Akamai offers its customers ways to have the edge servers ignore query strings, or even process many valid ones, these companies, according to Maggs, can easily forestall such an assault by using the right configuration when they hook into the network. But they just don't do that. "This attack doesn't happen in practice, so customers don't bother," he says.

Of course, if a miscreant were ever to take advantage of this weakness to paralyze a high-profile Web site, the companies using content-delivery networks to defend against denial-of-service attacks would probably pay more attention to all those pesky settings.

—David Schneider